arXiv:1404.5943v3 [cs.GT] 9 Jun 2014

Price of Anarchy for Auction Revenue*

Jason Hartline Darrell Hoy Sam Taggart

Northwestern University

June 10, 2014

Abstract

This paper develops tools for welfare and revenue analyses of Bayes-Nash equilibria in asymmetric auctions with single-dimensional agents. We employ these tools to derive *price of anarchy* results for social welfare and revenue. Our approach separates the standard smoothness framework [e.g., Syrgkanis and Tardos, 2013] into two distinct parts, isolating the analysis common to any auction from the analysis specific to a given auction. The first part relates a bidder's contribution to welfare in equilibrium to their contribution to welfare in the optimal auction using the price the bidder faces for additional allocation. Intuitively, either an agent's utility and hence contribution to welfare is high, or the price she has to pay for additional allocation is high relative to her value. We call this condition *value covering*; it holds in every Bayes-Nash equilibrium of any auction. The second part, *revenue covering*, relates the prices bidders face for additional allocation to the revenue of the auction, using an auction's rules and feasibility constraints. Combining the two parts gives approximation results to the optimal welfare, and, under the right conditions, the optimal revenue. In mechanisms with reserve prices, our welfare results show approximation with respect to the optimal mechanism with the same reserves.

As a center-piece result, we analyze the single-item first-price auction with individual monopoly reserves (the price that a monopolist would post to sell to that agent alone; these reserves are generally distinct for agents with values drawn from distinct distributions). When each distribution satisfies the regularity condition of Myerson [1981], the auction's revenue is at least a $2e/(e-1) \approx 3.16$ approximation to the revenue of the optimal auction. We also give bounds for matroid auctions with first-price or all-pay semantics, and the generalized first-price position auction. Finally, we give an extension theorem for simultaneous composition, i.e., when multiple auctions are run simultaneously, with single-valued, unit-demand agents.

*We thank Vasilis Syrgkanis for comments on a prior version of this paper for which simultaneous composition did not hold, suggesting the study of simultaneous composition and for perspective on price-of-anarchy methodology.

1 Introduction

The first step of a classical microeconomic analysis is to solve for equilibrium. Consequently, such analysis is restricted to settings for which equilibrium is analytically tractable; these settings are often disappointingly idealistic. Methods from the price of anarchy provide an alternative approach. Instead of solving for equilibrium, properties of equilibrium can be quantified from consequences of best response. These methods have been primarily employed for analyzing social welfare. While welfare is a fundamental economic objective, there are many other properties of economic systems that are important to understand. This paper gives methods for analyzing the price of anarchy for revenue.

Equilibrium requires that each agent's strategy be a best response to the strategies of others. A typical price-of-anarchy analysis obtains a bound on the social welfare (the sum of the revenue and all agent utilities) from a lower bound an agent's utility implied by best response. Notice that the agents themselves are each directly attempting to optimize a term in the objective. This property makes social welfare special among objectives. Can simple best-response arguments be used to quantify and compare other objectives? This paper considers the objective of revenue, i.e., the sum of the agent payments. Notice that each agent's payment appears negatively in her utility and, therefore, she prefers smaller payments; collectively the agents prefer smaller revenue.

The agenda of this paper parallels a recent trend in mechanism design. Mechanism design looks at identifying a mechanism with optimal performance in equilibrium. Optimal mechanisms tend to be complicated and impractical; consequently, a recent branch of mechanism design has looked at quantifying the loss between simple mechanisms and optimal mechanisms. These simple (designed) mechanisms have carefully constructed equilibria (typically, the truthtelling equilibrium). The restriction to truthtelling equilibrium, though convenient in theory, is problematic in practice [Ausubel and Milgrom, 2006]. In particular, this truthtelling equilibrium is specific to an ideal agent model and tends to be especially non-robust to out-of-model phenomena. The price of anarchy literature instead considers the analysis of the performance of simple mechanisms absent a carefully constructed equilibrium.

As an example, consider the single-item first-price auction, in which agents place sealed bids, the auctioneer selects the highest bidder to win, and the winner pays her bid. The fundamental tradeoff faced by the agents in selecting a bidding strategy is that higher bids correspond to higher chance of winning (which is good) but higher payments (which is bad). This first-price auction is the most fundamental auction in practice and it is the role of auction theory to understand its performance. When the agents' values for the item are drawn independently and identically, first-price equilibria are well-behaved: the symmetry of the setting enables the easy solving for equilibrium [Krishna, 2009], the equilibrium is unique [Chawla and Hartline, 2013; Lebrun, 2006; Maskin and Riley, 2003], and the highest valued agent always wins (i.e., the social welfare is maximized). When the agents' values are non-identically distributed, analytically solving for equilibrium is notoriously difficult. For example, Vickrey [1961] posed the question of solving for equilibrium with two agents with values drawn uniformly from distinct intervals; this problem was finally resolved half a century later by Kaplan and Zamir [2012].

Price-of-anarchy style analysis allows us to make general statements about behavior in equilibrium without needing an analytical characterization of equilibrium. For example, a recent analysis of Syrgkanis and Tardos [2013] shows that the first-price auction's social welfare in equilibrium is at least an $e/(e-1) \approx 1.58$ approximation to the optimal social welfare, and moreover, this bound continues to hold if multiple items are sold simultaneously by independent first-price

auctions. Importantly, this price-of-anarchy analysis sidesteps the intractability of solving for equilibrium and instead derives its bounds from simple best-response arguments.

1.1 Methods

Our analysis breaks down the problem of analyzing welfare and revenue into two parts. The first part, *value covering*, considers each agent individually and requires that an agent's contribution to welfare in BNE and her expected prices for additional allocation combine to approximate her contribution to welfare in the optimal mechanism. This analysis uses only properties of BNE. The second part, *revenue covering*, captures the relevant mechanism-specific details and considers the auction rules in aggregate across the agents. It requires that the auction's expected revenue approximately covers the expected prices for additional allocation across agents. The two parts combine to give price of anarchy bounds for welfare with and without reserves; adding the characterization of revenue in Bayes-Nash equilibrium of Myerson [1981] yields revenue approximation results.

Our analysis begins by translating the payments in any auction into *equivalent bids*: the first-price bids or payments if the payment rule of the mechanism used first-price semantics. This allows us to reduce the optimization problem a bidder faces into the same problem a bidder in the first price auction faces. From this viewpoint we show that in a Bayes-Nash equilibrium of *any auction* an agent's contribution to welfare and her expected prices for additional allocation combine to cover an $(e-1)/e$ fraction of her contribution to welfare in the optimal mechanism. Intuitively, either an agent's utility and hence contribution to welfare is high, or the price she has to pay for additional allocation is high relative to her value. If there the prices agents see for additional allocation correspond directly to the revenue of the mechanism, then combining across all agents gives that the welfare and revenue in Bayes-Nash equilibrium combine to approximate the welfare of the optimal mechanism. With reserve prices, considering only the agents with values above their reserve gives an approximation result to the optimal auction subject to the same reserves.

We then use the characterization of revenue in Bayes-Nash equilibrium of Myerson [1981] to repurpose our welfare analysis for revenue. The same covering condition that holds for bidders' values also holds for their (positive) virtual values: if a bidder has a positive virtual value, her contribution to virtual welfare in equilibrium and expected prices for additional allocation combine to cover an $(e-1)/e$ fraction of her contribution to virtual welfare in the optimal mechanism. If the revenue of the mechanism covers the the prices agents see for additional allocation, then the virtual welfare from positive virtual valued agents approximates the revenue of the revenue-optimal mechanism. If the impact of negative virtual-valued agents is small enough, equilibrium revenue will then approximate the optimal revenue. Setting monopoly reserve prices ensures this in a first-price auction. We also show that in first-price and all-pay auctions, if every bidder must compete with at least one other bidder with value drawn from her same distribution, then the revenue of a mechanism is at least half the virtual surplus from only positive virtual-valued agents.

1.2 Results

For single-item and single-dimensional matroid auctions (where the feasibility constraint is given by a matroid set system), we give welfare and revenue price of anarchy results with both first-price and all-pay payment semantics. The first-price variants of these auctions (a) solicit bids,

(b) choose an outcome to optimize the sum of the reported bids of served agents, and (c) charge the agents that are served their bids. These results are compatible with reserve prices. The all-pay variants of these auctions (a) solicit bids, (b) choose an outcome to optimize the sum of the reported bids of served agents, and (c) charge all agents their bids.

Welfare. In first-price auctions, we show that the price of anarchy for welfare is at most $2e/(e-1)$, with or without reserves. These results also extend to the generalized first-price position auction. For all-pay auctions in the above environments, we show the price of anarchy for welfare is $3e/(e-1)$. Tighter results are known with no reserves ($e/(e-1)$ and 2, respectively [Syrgkanis and Tardos, 2013]); the results with reserves are new.

Revenue. For first-price auctions with monopoly reserves in regular, single-parameter environments, we show that the price of anarchy for revenue is at most $2e/(e-1)$. The same bound holds in the generalized first-price position auction with monopoly reserves. If instead of reserves each bidder must compete with at least one duplicate bidder, the price of anarchy for revenue in first-price auctions is at most $3e/(e-1)$; in all-pay auctions, at most $4e/(e-1)$.

Simultaneous Composition. We also show via an extension theorem that the above bounds hold when auctions are run simultaneously if agents are *unit-demand* and *single-valued* across the outcomes of the auctions.

1.3 Related Work

Understanding welfare in games without solving for equilibrium is a central theme in the smooth games framework of Roughgarden [2009] and the smooth mechanisms extension of Syrgkanis and Tardos [2013]. A core principle of smoothness is that the precise manner or assumptions used in proving the smoothness property dictate how broadly the result extends. One view of our work is that we limit our proofs in just the right way that allows for extensions to revenue approximations.

Our framework refines the smoothness framework for Bayesian games in three notable ways. First, we decompose smoothness into two components, separating the specifics of a mechanism from the actions of a best-responding agent in any auction equilibrium. Second, because we focus on the optimization problem that individual bidders are facing, we can attain results that only hold for certain bidders — for instance, bidders with values above their reserve prices. Third, we only consider the Bayesian setting, which allows us to use the BNE characterization of Myerson [1981] to approximate revenue and relate other auctions to the first-price auction via equivalent bids.

A number of papers looking at revenue guarantees for the welfare-optimal Vickrey-Clarke-Groves (VCG) mechanism in asymmetric settings. Hartline and Roughgarden [2009] show that VCG with monopoly reserves or duplicate bidders achieves revenue that is a constant approximation to the revenue optimal auction. Dhangwatnotai et al. [2010] show that the single-sample mechanism, essentially VCG using a single bid as a reserve, achieves approximately optimal revenue in broader settings. Roughgarden et al. [2012] showed that in broader environments, including matching settings, limiting the supply of items in relation to the number of bidders gives a constant approximation to the optimal auction.

In the economics literature, a number of papers have explored properties of asymmetric first-price auctions. Kirkegaard [2009] shows that understanding the ratios of expected payoffs in equilibrium can be easier than understanding equilibrium and lead to insights about equilibria. Kirkegaard [2012] shows that some properties of distributions can be used to compare revenue of the first price auction to revenue of the second price auction. Lebrun [2006] and

Maskin and Riley [2003] establish equilibrium uniqueness in the asymmetric setting with some assumptions on the distributions of agents.

2 Preliminaries

Bayesian Mechanisms. This paper considers mechanisms for n single-dimensional agents with linear utility. Each agent has a private value for service, v_i, drawn independently from a distribution F_i over valuation space V_i. We write $\mathbf{F} = \prod_i F_i$ and $\mathbf{V} = \prod_i V_i$ to denote the joint value distribution and space of valuation profiles, respectively. A *mechanism* consists of a bid allocation rule $\tilde{\mathbf{x}}$ and a payment rule $\tilde{\mathbf{p}}$, mapping actions of agents to allocations and payments respectively. Each agent i draws their private value v_i from F_i and selects an action according to some strategy $s_i : V_i \to A_i$, where A_i is the set of possible actions for i. We write $\mathbf{s} = (s_1, \ldots, s_n)$ to denote the vector of agents' strategies. Given the actions $\mathbf{a} = (a_1, \ldots, a_n)$ selected by each agent, the mechanism computes $\tilde{\mathbf{x}}(\mathbf{a})$ and $\tilde{\mathbf{p}}(\mathbf{a})$. Each agent's utility is $\tilde{u}_i(\mathbf{a}) = v_i\tilde{x}_i(\mathbf{a}) - \tilde{p}_i(\mathbf{a})$.

Mechanisms typically operate with constraints on permissible allocations. A *feasibility environment* specifies the set of feasible allocation vectors. Mechanisms for a feasibility environment must choose only allocations from the feasible set. The simplest example is a single-item auction, in which at most one person at a time can be served. This paper assumes feasibility environments are *downward-closed*: if $(x_1, \ldots, x_k, \ldots, x_n)$ is feasible, so is $(x_1, \ldots, 0, \ldots, x_n)$. We will often consider the special case of *matroid* environments, in which the set of feasible allocations correspond to the independent sets of a matroid set system.

Given a strategy profile $\mathbf{s}$, we often consider the expected allocation and payment an agent faces from choosing some action $a_i \in A_i$, with expectation taken with respect to other agents' values and actions induced by $\mathbf{s}$. We treat $\mathbf{s}$ as implicit and write $\tilde{x}_i(a_i) = \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{x}_i(a_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$, with $\tilde{p}_i(a_i)$ and $\tilde{u}_i(a_i)$ defined analogously. Given $\mathbf{s}$ implicitly, we also consider values as inducing payments and allocations. We write $\mathbf{x}(\mathbf{v}) = \tilde{\mathbf{x}}(\mathbf{s}(\mathbf{v}))$ and $\mathbf{p}(\mathbf{v}) = \tilde{\mathbf{p}}(\mathbf{s}(\mathbf{v}))$, respectively. Furthermore, we can denote agent i's interim allocation probability and payment by $x_i(v_i) = \tilde{x}_i(s_i(v_i))$ and $p_i(v_i) = \tilde{p}_i(s_i(v_i))$. We define $u(\mathbf{v})$ and $u_i(v_i)$ similarly. In general, we use a tilde to denote outcomes induced by actions, and omit the tilde when indicating outcomes induced by values. We refer to $\tilde{\mathbf{x}}$ as the *bid allocation rule*, to distinguish it from $\mathbf{x}$, the *allocation rule*. We adopt a similar convention with other notation.

Bayes-Nash Equilibrium. A strategy profile $\mathbf{s}$ is in *Bayes-Nash equilibrium* (BNE) if for all agents i, $s_i(v_i)$ maximizes i's interim utility, taken in expectation with respect to other agents' value distributions $\mathbf{F}_{-i}$ and their actions induced by $\mathbf{s}$. That is, for all i, v_i, and alternative actions a': $\mathbf{E}_{\mathbf{v}_{-i}}[\tilde{u}_i(\mathbf{s}(\mathbf{v}))] \geq \mathbf{E}_{\mathbf{v}_{-i}}[\tilde{u}_i(a', \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$.

We will consider only mechanisms where agents can gain from participation, regardless of their value — that is, mechanisms that are *interim individually rational*. We will thus assume that in any auction every bidder has at least one withdraw action w_i such that $\tilde{x}_i(w_i, \mathbf{a}_{-i}) = 0$ and $\tilde{p}_i(w_i, \mathbf{a}_{-i}) = 0$ for any $\mathbf{a}_{-i}$. We assume all mechanisms have at least one such action for each agent. In any BNE, each agent has the option to withdraw and must therefore get non-negative utility.

Myerson [1981] characterizes the interim allocation and payment rules that arise in BNE. These results are summarized in the following theorem.

Theorem 1 (Myerson, 1981)**.** *For any mechanism and value distribution* $\mathbf{F}$ *in BNE,*

1. *(monotonicity) The interim allocation rule* $x_i(v_i)$ *for each agent is monotone non-decreasing in* v_i*.*

2. *(payment identity) The interim payment rule satisfies* $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz$.
3. *(revenue equivalence) Mechanisms and equilibria which result in the same interim allocation rule* $\mathbf{x}(\mathbf{v})$ *must therefore have the same interim payments as well.*

Mechanism Design Objectives. We consider the problem of maximizing two primary objectives in expectation at BNE: welfare and revenue. The revenue of a mechanism M is the total payment of all agents. Mechanism M's expected revenue for $\mathbf{v} \sim \mathbf{F}$ in a given Bayes-Nash equilibrium $\mathbf{s}$ is denoted $\textsc{Rev}(M) = E_{\mathbf{v}}[\sum_i p_i(\mathbf{v})]$. The welfare of a mechanism M is the total utility of all participants including the auctioneer; denoted $\textsc{Welfare}(M) = \textsc{Rev}(M) + E_{\mathbf{v}}[\sum_i u_i(\mathbf{v})] = E_{\mathbf{v}}[v_i x_i(\mathbf{v})]$. We will use the term surplus synonymously with welfare.

Our welfare benchmark is the mechanism that always serves the highest valued feasible agents. That is, we seek to approximate $\textsc{Welfare}(\textsc{Opt}) = \mathbf{E}_{\mathbf{v}}[\max_{\mathbf{x}^*} \sum_i v_i x_i^*]$. This can be implemented via the Vickrey-Clarke-Groves (VCG) mechanism. We measure a mechanism M's welfare performance by the *Bayesian price of anarchy for welfare*, given by

$$\max_{\mathbf{F}, \mathbf{s} \in \text{BNE}(M,\mathbf{F})} \frac{\textsc{Welfare}(\textsc{Opt})}{\textsc{Welfare}(M)},$$

where $\text{BNE}(M, \mathbf{F})$ is the set of BNE for M under value distribution $\mathbf{F}$.

For revenue, we will make strong use of the characterization of revenue in Myerson [1981] that follows from Theorem 1:

Lemma 2. *In BNE, the ex ante expected payment of an agent is* $\mathbf{E}_{v_i}[p_i(v_i)] = \mathbf{E}_{v_i}[\phi_i(v_i)x_i(v_i)]$, *where* $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ *is the* virtual value *for value* v_i. *It follows that* $\textsc{Rev}(M) = E_{\mathbf{v}}[\sum_i p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\sum_i \phi_i(v_i)x_i(\mathbf{v})]$.

Using this result, Myerson [1981] derives the revenue-optimal mechanism for any value distribution $\mathbf{F}$. This mechanism is parameterized by the value distribution $\mathbf{F}$, and the optimality is in expectation over $\mathbf{v} \sim \mathbf{F}$. We specifically consider distributions with no point masses where $\phi_i(v_i)$ is monotone in v_i for each i. Such distributions are said to be *regular*. If each agent has a regular distribution, then the revenue-optimal mechanism selects the allocation which maximizes $\sum_i \phi_i(v_i)x_i(\mathbf{v})$. We will seek to minimize the *Bayesian price of anarchy for revenue*,

$$\max_{\mathbf{F} \in \mathcal{R}, \mathbf{s} \in \text{BNE}(M,\mathbf{F})} \frac{\textsc{Rev}(\textsc{Opt}_{\mathbf{F}})}{\textsc{Rev}(M)},$$

where $\mathcal{R}$ is the set of regular distributions and $\textsc{Opt}_{\mathbf{F}}$ is the Bayesian revenue-optimal mechanism for value distribution $\mathbf{F}$.

3 Single-Item First Price Auction with Reserves

We begin by analyzing the single-item first price auction with per-bidder reserves, and show that it approximates the welfare of the welfare-optimal mechanism with the same reserves. With reserves of zero, this implies that the welfare of the first-price auction with no reserves approximates the welfare of the welfare-optimal auction. We then apply these techniques to revenue approximation, getting the result that the first-price auction with monopoly reserves approximates the revenue of the revenue-optimal mechanism.

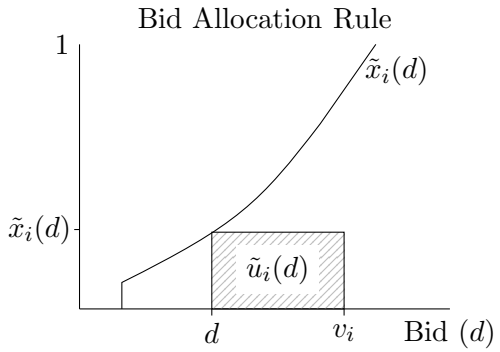

Figure 1: For any bid d, the area of a rectangle between $(d, \tilde{x}_i(d))$ and $(v_i, 0)$ on the bid allocation rule is the expected utility $\tilde{u}_i(d)$. The BNE bid b_i is chosen to maximize this area.

3.1 Welfare

We now show that the welfare and revenue of the first price auction together approximate the welfare of the welfare-optimal auction, specifically:

$$\textsc{Welfare}(\text{FPA}_{\mathbf{r}}) + \textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \tfrac{e-1}{e}\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}}). \tag{1}$$

Our proof proceeds by first analyzing the optimization problem of a bidder, then relating that optimization problem to welfare and revenue. We conclude with the following theorem.

Theorem 3. *The welfare in any BNE of the first price auction with reserves* $\mathbf{r}$ *is at least a* $\frac{2e}{e-1}$*-approximation to the welfare of the welfare-optimal mechanism that serves no agent with* $v_i < r_i$*.*

Note that equation (1) resembles closely the inequality in the smooth games and mechanism frameworks [Syrgkanis and Tardos, 2013; Roughgarden, 2009]. It differs primarily in that we are not defining a specific deviation, but rather deriving bounds explicitly from BNE.

A bidder's optimization problem. Consider the optimization problem faced by a bidder i with value v_i in the first price auction. Bidder i's expected utility for a possible bid d is $\tilde{u}_i(d) = (v_i - d)\tilde{x}_i(d)$, where $\tilde{x}_i(d)$ is the interim bid allocation rule faced by the bidder. Let b_i be her best response bid given her value v_i. That is, b_i maximizes $\tilde{u}_i(d)$. If we plot the bid allocation rule $\tilde{x}_i(d)$ for any alternate bid d, then $\tilde{u}_i(b_i)$ is precisely the area of the rectangle in the lower right; see Figure 1.

When other bidders have realized values and submitted bids, bidder i wins only if her bid exceeds both her reserve and the bids of other players. Consequently, $\tau_i(\mathbf{v}_{-i}) = \max(r_i, \mathbf{b}_{-i}(\mathbf{v}_{-i}))$ is the price a bidder must pay to win; we will formally refer to it as her *threshold bid*. As we are in the Bayesian setting, a bidder is not reacting to this threshold, but is acting in expectation over the types and actions of her competitors. These actions induce a distribution over threshold bids. The cumulative distribution function of threshold bids for a bidder i is precisely the bid allocation rule $\tilde{x}_i$.

We will also refer to thresholds using the probability of allocation that they represent achieving. Let $\tau_i(x)$ refer to the smallest bid that achieves allocation of at least x, hence $\tau_i(x) = \min\{\ b \mid \tilde{x}_i(b) \geq x\}$. Note that $\tau_i(x)$ is the price an agent faces to attain allocation x.

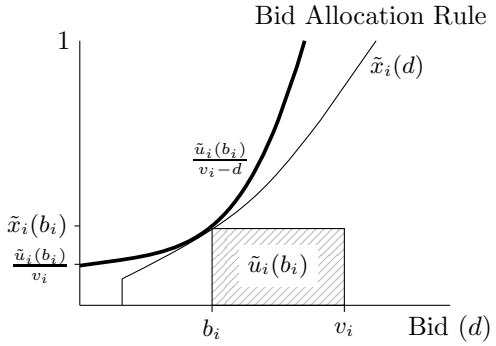

(a) As b_i is a best-response to the actions of other agents, the indifference curve $\tilde{u}_i(b_i)/(v_i - d)$ upper bounds $\tilde{x}_i(d)$.

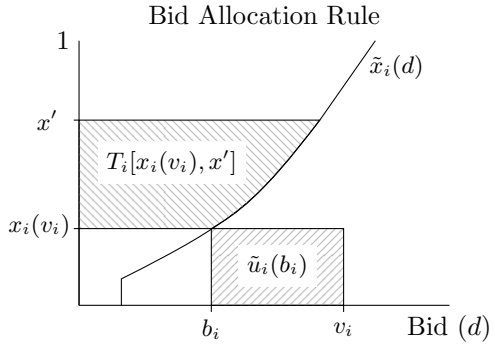

(b) The cumulative additional threshold bid $T_i[x_i(v_i), x']$ prevents bidders from bidding to receive allocation x'.

Figure 2

Let $B_{-i}(b)$ be the cumulative distribution function of the highest bids from other bidder. Then $\tau_i(x)$ is either the reserve r_i or the bid required to beat the highest bid from other agents a x fraction of the time, $\tau_i(x) = \max(r_i, B_{-i}^{-1}(x))$.[1]

For an alternate allocation probability x', the *cumulative additional threshold* $T_i[x_i, x'] = \int_{x_i}^{x'} \tau_i(z)\, dz$ as a measure of the prices an agent faces to attain additional allocation between x_i and x'. This is illustrated in Figure 2b. The cumulative additional threshold is the quantity we will be using to relate welfare contributions in equilibrium and the optimal allocation.

Relating Contributions to First-Price and Optimal Welfare: Let $\mathbf{x}^*$ be the allocation rule from the welfare optimal mechanism that serves no agent with value $v_i < r_i$, $\textsc{Opt}_{\mathbf{r}}$. Thus $\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}}) = \sum_i \mathbf{E}_{v_i}[v_i x_i^*(v_i)]$, where $v_i x_i^*(v_i)$ is bidder i's contribution to the optimal welfare. We will now approximate each bidder's contribution to the optimal welfare individually, using the bidder's contribution to welfare in the first-price auction, i.e., $v_i x_i(v_i)$, and a fraction of the revenue in the first-price auction.

The proof of Theorem 3 proceeds in two steps:

1. *Value Covering*: A bidder's contribution to welfare in the $\text{FPA}_{\mathbf{r}}$ and *cumulative additional threshold* together approximate her contribution to welfare in any alternate allocation. (Lemma 4)
2. *Revenue Covering*: The revenue of the $\text{FPA}_{\mathbf{r}}$ approximates the cumulative additional threshold of all agents. (Lemma 5)

The final approximation result follows by summing the value covering inequality across agents, taking expectation over values, and combining with revenue covering. When value covering is employed to show an approximation to the welfare induced by an allocation rule x_i^*, the alternate allocation x' used for every bidder and value will be precisely $x' = x_i^*(v_i)$.

Lemma 4 (Value Covering)**.** *In any BNE of* $\text{FPA}_{\mathbf{r}}$*, for any bidder i with value v_i and alternate feasible allocation x',*

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e} v_i x'. \tag{2}$$

[1] If B_{-i} is not invertible, then define $B_{-i}^{-1}(x)$ to be the function $B_{-i}^{-1}(x) = \inf\{b \mid \tilde{x}_i(b) \geq x\}$.

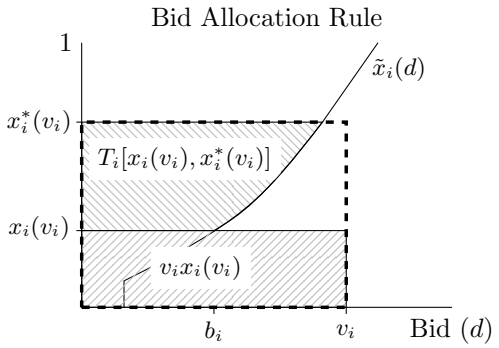

(a) Lemma 4 shows the shaded areas cover a $(e-1)/e$ fraction of the dashed box, bidder i's contribution to the optimal welfare ($v_i x_i^*(v_i)$).

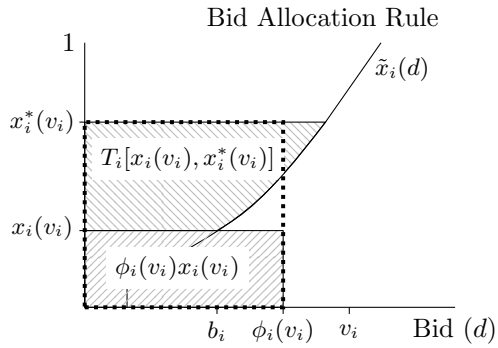

(b) Lemma 7 shows the shaded areas cover an $(e-1)/e$ fraction of i's contribution to the optimal virtual welfare, $\phi_i(v_i)x_i^*(v_i)$.

Figure 3

Proof. We will prove value covering in two steps: first, by developing a lower bound $\underline{T}$ on the cumulative additional threshold T; second, by optimizing to get the worst such bound. The proof can also be done using a modification of the first-price bid deviation approach of Syrgkanis and Tardos [2013].

Lowerbounding T. In best responding, bidder i chooses an action which maximizes her utility. If b_i is a best response bid, then for any alternate bid d, $\tilde{u}_i(b_i) \geq (v_i - d)\tilde{x}_i(d)$, hence $\tilde{x}_i(d) \leq \frac{\tilde{u}_i(b_i)}{v_i - d}$. With equality, this bound gives an indifference curve for bidder i; it is the alternate bid allocation rule that would lead to her being indifferent over all reasonable bids (see Figure 2a). Call $\underline{T}_i[x_i(v_i), x']$ the expected threshold bid from the indifference curve, then $\underline{T}_i[x_i(v_i), x'] = \int_{x_i(v_i)}^{x'} \max(0, v_i - u_i(v_i)/z)\ dz$ and

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq v_i x_i(v_i) + \underline{T}_i[x_i(v_i), x'] \geq u_i(v_i) + \underline{T}_i[u_i(v_i)/v_i, x']. \tag{3}$$

The last inequality follows because the threshold bid for i when i is winning must be less than her bid, thus the cumulative threshold below her allocation is upperbounded by her expected payment, hence $b_i(v_i)x_i(v_i) \geq T_i[0, x_i(v_i)] \geq \underline{T}_i[u_i(v_i)/v_i, x_i(v_i)]$.

Worst-case $\underline{T}_i$. Evaluating the integral for $\underline{T}_i[u_i(v_i)/v_i, v_i]$ gives $u_i(v_i) + \underline{T}_i[u_i(v_i)/v_i, x'] = v_i x' + u_i(v_i) \ln \frac{u_i(v_i)}{v_i x'}$. Holding $v_i x'$ fixed and minimizing with respect to $u_i(v_i)$ yields a minimum at $u_i(v_i) = \frac{v_i x'}{e}$, hence

$$u_i(v_i) + \underline{T}_i[u_i(v_i)/v_i, x'] \geq \tfrac{e-1}{e} v_i x'. \tag{4}$$

Combining (3) and (4) gives exactly our desired result, (2). □

We now show that in the first price auction, the expected revenue is greater than the cumulative additional threshold necessary to achieve any alternate feasible allocation $\mathbf{x}'$, which we can then combine with value covering to give a welfare approximation result. While value covering depended critically on equilibrium (or at least on bidders best responding), revenue covering will only depend on the form of the first price auction, and will thus hold for arbitrary

(not necessarily BNE) bidding strategies that satisfy a light participation requirement (that is always satisfied in BNE). We call a bidding strategy *participatory* if every bidder i bids above their reserve as long as there is some bid in $[r_i, v_i]$ that would give them strictly positive utility.

Lemma 5 (Revenue Covering)**.** *For any participatory bidding strategies* $\mathbf{s}$ *and alternative allocation* $\mathbf{x}'$*,*

$$\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \sum_{i\,:\,v_i > r_i} T_i[\tilde{x}_i(s_i(v_i)), x_i']. \tag{5}$$

Proof. It suffices to show that for bidder i with value above her reserve,

$$\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq T_i[\tilde{x}_i(s_i(v_i)), 1]. \tag{6}$$

From equation (6), multiplying by x_i', summing over agents, and observing that $T_i[\tilde{x}_i(s_i(v_i)), z]$ is convex in z concludes the lemma:

$$\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \sum\nolimits_i x_i' T_i[\tilde{x}_i(s_i(v_i)), 1] \geq \sum\nolimits_i T_i[\tilde{x}_i(s_i(v_i)), x_i']. \tag{7}$$

We now show equation (6). By the participatory assumption, a bidder i with value above her reserve bids above her reserve r_i. Thus the cumulative additional threshold when playing $s_i(v_i)$ is bounded by the cumulative additional threshold when bidding the reserve, hence $T_i[\tilde{x}_i(s_i(v_i)), 1] \leq T_i[\tilde{x}_i(r_i), 1]$. Then, using the definition of T and the monotonicity of $B_{-i}^{-1}(x)$, we have:

$$T_i[\tilde{x}_i(s_i(v_i)), 1] \leq T_i[\tilde{x}_i(r_i), 1] = \int_{\tilde{x}_i(r_i)}^{1} \max(r_i, B_{-i}^{-1}(z))\ dz = \int_{\tilde{x}_i(r_i)}^{1} B_{-i}^{-1}(z)\ dz. \tag{8}$$

As the revenue of a first price auction is the expected highest bid, and $B_{-i}^{-1}(z)$ is the inverse of the cumulative distribution function of highest bid from bidders aside from i, $\int_0^1 B_{-i}^{-1}(z)\ dz = \mathbf{E}_{\mathbf{v}}[\max_{j\neq i,\ s_i(v_j)\geq r_j} s_i(v_j)] \leq \textsc{Rev}(\text{FPA}_{\mathbf{r}})$. Combining this with monotonicity of $B_{-i}^{-1}(z)$ gives

$$T_i[\tilde{x}_i(s_i(v_i)), 1] \leq \int_{\tilde{x}_i(r_i)}^{1} B_{-i}^{-1}(z)\ dz \leq \int_{0}^{1} B_{-i}^{-1}(z)\ dz \leq \textsc{Rev}(\text{FPA}_{\mathbf{r}}). \tag{9}$$

Chaining (8) and (9) gives $\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq T_i[\tilde{x}_i(s_i(v_i)), 1]$. □

We now combine value and revenue covering to attain an approximation to the optimal welfare.

Proof of Theorem 3. We apply value and revenue covering with $\mathbf{x}' = \mathbf{x}^*(\mathbf{v})$ for all agents with values above their reserves. As we assume to have no point masses in our distributions, the contribution to optimal welfare from bidders with values exactly their reserves is 0. Taking equation (2) summed over all agents, in expectation over values above reserves and combining with (5) gives $\textsc{Welfare}(\text{FPA}_{\mathbf{r}}) + \textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \frac{e-1}{e}\textsc{Welfare}(\textsc{Opt}_{\mathbf{r}})$. As $\textsc{Welfare}(\text{FPA}_{\mathbf{r}}) \geq \textsc{Rev}(\text{FPA}_{\mathbf{r}})$, $\textsc{Welfare}(\text{FPA}_{\mathbf{r}})$ is then a $2e/(e-1)$ approximation to $\textsc{Opt}_{\mathbf{r}}$. □

The following are the main ideas and differences between the proof above and the proof of Syrgkanis and Tardos [2013] that enables treatment of reserve prices. The Syrgkanis and Tardos [2013] result can be viewed as combining value covering and revenue covering in one step (via the smoothness definition). Their equation has bidders' utilities where we have the bidders'

surpluses and they have the full cumulative threshold where we have the cumulative additional threshold. The thresholds that a bidder faces that correspond to bids of other bidders translate to revenue and can be thus bounded by a revenue covering argument. Reserve prices, however, induce thresholds that do not correspond to bids of other bidders. A participatory bidder, however, will bid above the reserve when her value is above the reserve. Therefore, the bidder's payment will always compensate for the part of the threshold distribution that corresponds to the reserve price. Because we use surplus (utility and payment) instead of utility to account for the threshold lost to the reserves, our analysis loses a factor of two on the no-reserves bound of Syrgkanis and Tardos [2013].

3.2 Revenue

In the tradition of Bayesian mechanism design, we will prove a revenue approximation result by reducing to the welfare approximation above. Let x_i^* now denote the allocation rule from the revenue optimal auction of Myerson [1981]. Recall that for welfare, we approximated each bidders contribution to the optimal welfare using the cumulative additional threshold and contribution to equilibrium welfare. For revenue, we will instead approximate each bidders contribution to the optimal *virtual* welfare using the cumulative additional threshold and contribution to *virtual* welfare in equilibrium. This approximation will hold for agents with positive virtual values. To mitigate the impact of agents with negative virtual values, we set monopoly reserves at $r_i = \phi_i^{-1}(0)$, which by our assumption of regular distributions with no point masses eliminates all agents with negative virtual values and no agents with positive virtual values. Combining the two will yield an approximation result to the revenue optimal auction.

Theorem 6. *The revenue in any BNE of the first price auction with monopoly reserves (*$\text{FPA}_{\mathbf{r}}$*) and agents with regularly distributed values is at least a* $2e/(e-1)$*-approximation to revenue of the optimal auction.*

We will begin by showing the analogue of value covering, *virtual* value covering, in which a bidder's contribution to equilibrium virtual welfare and cumulative additional threshold combine to approximate her contribution to the optimal virtual welfare.

Lemma 7 (Virtual Value Covering)**.** *In any BNE of* $\text{FPA}_{\mathbf{r}}$*, for any bidder* i *with value* v_i *such that* $\phi_i(v_i) \geq 0$*, and alternate allocation* x'*,*

$$\phi_i(v_i)x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e}\phi_i(v_i)x'. \tag{10}$$

Proof. This follows directly from value covering (Lemma 4) — see Figure 3b for an illustration. Since virtual values are smaller than values, then the cumulative additional threshold is larger as a fraction of virtual surplus than surplus, $\frac{T_i[x_i(v_i),x']}{\phi_i(v_i)x'} \geq \frac{T_i[x_i(v_i),x']}{v_i(v_i)x'}$, while the virtual welfare in equilibrium as a fraction of the virtual welfare in the alternate allocation will be the same as the welfare in equilibrium as a fraction of welfare in the alternate allocation. Combining with the value covering condition (2) then gives:

$$\tfrac{1}{\phi_i(v_i)x'}\left(\phi_i(v_i)x_i(v_i) + T_i[x_i(v_i), x']\right) \geq \tfrac{1}{v_i x'}\left(v_i x_i(v_i) + T_i[x_i(v_i), x']\right) \geq \tfrac{e-1}{e}.$$

Multiplying through by $\phi_i(v_i)x'$ gives our desired result. ∎

Proof of Theorem 6. We apply virtual value covering and revenue covering with $\mathbf{x}' = \mathbf{x}^*(\mathbf{v})$. As no agent with negative virtual value is served in the first-price auction with monopoly reserves or the optimal auction, equation (10) holds for all bidders, not only those with positive virtual values. Summing over all players with $\phi_i(v_i) \geq 0$ and taking expectation over values gives

$$\begin{aligned}\sum_i \mathbf{E}_{v_i}\left[T_i[x_i, x_i^*(v_i)]\right] + \sum_i \mathbf{E}_{v_i}\left[\phi_i(v_i)x_i(v_i)\right] &\geq \tfrac{e-1}{e}\sum_i \mathbf{E}_{v_i}\left[\phi_i(v_i)x_i^*(v_i)\right] \\ &= \tfrac{e-1}{e}\textsc{Rev}(\textsc{Opt}) \qquad (11)\end{aligned}$$

Each term on the left side is at most the first-price revenue, by revenue covering and Lemma 2, so

$$2\textsc{Rev}(\text{FPA}_{\mathbf{r}}) \geq \tfrac{e-1}{e}\textsc{Rev}(\textsc{Opt}). \qquad (12)$$

Thus $\textsc{Rev}(\text{FPA}_{\mathbf{r}})$ is at least a $\frac{2e}{e-1}$ approximation to $\textsc{Rev}(\textsc{Opt})$. □

3.2.1 Welfare and Revenue Lower Bounds

The approximation results we have given in this section for the single-item first-price auction are not known to be tight. For welfare with no reserves, the price of anarchy can be as bad as 1.15; we give such an example in Appendix A. Note the large gap between this lower bound and the best known upper bound of $\frac{e}{e-1} \approx 1.58$ [Syrgkanis and Tardos, 2013].

For revenue, the approximation ratio can be at least as bad as 2, using the same lower bound Hartline and Roughgarden [2009] show for VCG with monopoly reserves. The example has two bidders, one with deterministic value 1, the other with value drawn according to the equal revenue distribution with support over $[1, H]$ for some large H. With a light perturbation of the distribution the monopoly price for the second bidder is 1. Assuming ties go to bidder 2, an equilibrium exists where both players bid 1, giving revenue of 1. The optimal auction however can set a reserve of H for the second bidder and sell to the first bidder at price 1 if the reserve is not met, achieving a revenue of 2 as H grows.

4 Framework

In equilibria of the single-item first-price auction, we observed that agents with low expected utility had high expected threshold bids. Because high thresholds were connected to high payments, we could conclude that the first price auction is both approximately welfare- and revenue-optimal. The goal for this section is to build up a framework for making this same argument for mechanisms with different payment semantics, such as all-pay auctions. In particular, we seek to prove results about behavior in Bayes-Nash equilibrium while *ignoring* the particular payment semantics of each auction. We begin by defining *equivalent bids*, which allow us to reduce the optimization problem a bidder faces in any auction to the problem faced in a first-price auction. This will allow us to reduce much of the analysis of general auctions for single-dimensional agents to the single-item first-price auction analysis of Section 3.

4.1 Equivalent Bids

Utility-maximizing agents must balance two goals: getting allocated frequently, and getting allocated cheaply. In a first-price auction, agents bid to explicitly specify the tradeoff they are willing to make: their bid is the price they pay per unit of allocation. In general mechanisms,

for any agent i and any action a_i, define the *equivalent bid* for an action a_i to be $\beta_i(a_i) = \tilde{p}_i(a_i)/\tilde{x}_i(a_i)$; this can be thought of as the price per unit of allocation for that action. For first-price auctions, this is exactly the bid. For mechanisms with different payment semantics, $\beta_i(a_i)$ can still be thought of as an equivalent first-price bid for action a_i.

Equivalent Threshold Bid. In proving Theorem 3, we noted that b is the minimum payment necessary to get the allocation probability $\tilde{x}_i(b)$. We used this property to bound the distribution of other agents' bids. For auctions where this relationship is less clear, we think of agents partitioning the actions in their choice set by interim allocation probability, then for each probability consider only the cheapest such action in terms of price per unit of allocation. For each allocation probability z, let $\alpha_i(z)$ be that cheapest action and let the *equivalent threshold bid* $\tau_i(z) = \beta_i(\alpha_i(z))$ be the equivalent bid of the cheapest action. Formally, $\tau_i(z) = \min_{a_i:\tilde{x}_i(a_i)\geq z} \beta_i(a_i)$, with $\alpha_i(z)$ the arg min. Note that $\tau_i(z)$ depends on $\mathbf{s}$; for notational convenience, we suppress the strategy profile as an argument.

Cumulative Equivalent Threshold Bid. We can now use $\tau_i(z)$ to track the expense an agent faces from increasing their allocation. Specifically, assume an agent is playing some action a_i and seeks to increase their allocation probability to x'. The barrier to i doing so is the set of equivalent threshold bids in $[\tilde{x}_i(a_i), x']$. We can use this notion to measure i's expense for additional allocation. Define the *cumulative equivalent threshold bid* as $T_i[\tilde{x}_i(a_i), x'] = \int_{\tilde{x}_i(a_i)}^{x'} \tau_i(z)\, dz$. If $x' \leq \tilde{x}_i(a_i)$, then define $T_i[\tilde{x}_i(a_i), x'] = 0$. This quantity will function identically to its counterpart in Section 3, trading off against i's surplus as in Lemma 4, and translating into revenue as in Lemma 5. Note that because $\tau_i(z)$ is nondecreasing in z, $T_i[\tilde{x}_i(a_i), x']$ is convex in x'.

4.2 Covering Conditions and the Price of Anarchy

We now use equivalent bids and thresholds in place of first-price bids and thresholds to develop general analogues of the value and revenue covering conditions of Section 3.

Lemma 8 (Value Covering)**.** *Consider a mechanism M in BNE with induced allocation and payment rules $(\mathbf{x}, \mathbf{p})$, and an agent i with value v_i. For any $x' \in [0,1]$,*

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e} v_i x'. \tag{13}$$

The proof can now be done by reduction to the single-item first-price auction (Lemma 4) because bidders now face effectively the same optimization problem as in a single-item first-price auction. The proof is included in Appendix B.

To prove an approximation result for welfare or revenue, the *only* mechanism-specific detail which remains is specifying the relationship between T_i and the mechanism's revenue. Intuitively, we saw in Section 3 that if there is a relationship between revenue and the difficulty an agent faces in increasing their allocation once they have chosen to participate in the mechanism, then value covering allows us to show a welfare bound. To make this relationship concrete, we extend the definition of Lemma 5.

Definition 9 (Revenue Covering)**.** *A mechanism M is* μ-revenue covered *if for any (implicit) strategy profile $\mathbf{s}$, feasible allocation $\mathbf{x}'$, and action profile $\mathbf{a}$,*

$$\mu \textsc{Rev}(M) \geq \sum\nolimits_i T_i[\tilde{x}_i(a_i), x_i'].$$

Note that Definition 9 makes no mention of BNE. It must hold for any strategy profile. This is a stronger condition than Lemma 5, as it is not restricted to bidders with values above a set of reserves or bidders playing only participatory strategies.

As we already saw, revenue covering has a number of important consequences. First is a welfare bound.

Theorem 10. *If a mechanism is μ-revenue covered, then in any BNE it is a $(1+\mu)\frac{e}{e-1}-$approximation to the welfare of the optimal mechanism.*

Proof. Let $\mathbf{x}^*$ be the welfare-optimal allocation rule. For any value profile $\mathbf{v}$, Lemma 8 with $x' = x_i^*(\mathbf{v})$ yields that for each $\mathbf{x}$,

$$v_i x_i(v_i) + T_i[x_i(v_i), x_i^*(\mathbf{v})] \geq \tfrac{e-1}{e} v_i x_i^*(\mathbf{v}).$$

Summing over agents and using revenue covering gives $\sum_i v_i x_i(v_i) + \mu \textsc{Rev}(M) \geq \frac{e-1}{e} \sum_i v_i x_i^*(\mathbf{v})$. Taking expectation with respect to $\mathbf{v}$ and using $\textsc{Welfare}(M) \geq \textsc{Rev}(M)$, we get

$$(1+\mu)\frac{e}{e-1}\textsc{Rev}(M) \geq \frac{e}{e-1}\textsc{Welfare}(\textsc{Opt}). \tag{14}$$

□

4.3 Revenue Covering With Reserves

Not all agents need their thresholds covered for a welfare and revenue approximation result. For instance, in Section 3, we considered the first-price auction with monopoly reserves. Bidders with values below their reserve experience a threshold created by the reserve of the auction which does not translate into revenue. We showed instead that revenue covers the thresholds for bidders with values (and hence equilibrium bids) above their reserves, which was sufficient for approximation results because the optimal auction served no bidder with value below their reserve.

For the actions of such agents to be revenue covered, we do need bidders with values above their reserve to bid at least the reserve, just as in Section 3. Given a strategy profile $\mathbf{s}$ and value profile $\mathbf{v}$, define an action a_i to be participatory for $\mathbf{s}$ and $\mathbf{v}$ if the equivalent bid $\beta_i(a_i) \leq v_i$ and either $\tilde{x}_i(a_i) > 0$ or there is no alternate action a_i' that gives i strictly positive utility. Note that participatory is a much weaker assumption than equilibrium or best-response; it only specifies that bidders play some action that gives positive allocation if there is such an (IR) action. Thus in BNE, all bidders will play participatory actions.

Definition 11 (Revenue Covering with Reserves)**.** *A mechanism M is μ-revenue covered with reserves $\mathbf{r} = (r_1, \ldots, r_n)$ if for any (implicit) strategy profile $\mathbf{s}$, value profile $\mathbf{v}$, feasible allocation $\mathbf{x}'$, and participatory action profile $\mathbf{a}$,*

$$\mu \textsc{Rev}(M) \geq \sum\nolimits_{i: v_i > r_i} T_i[\tilde{x}_i(a_i), x_i'].$$

Note that revenue covering above reserves is a weaker condition than general revenue covering, and is not relying on any specific property of the mechanism, like whether it does or does not have reserves. For instance, the first price auction with no reserves will satisfy revenue covering with any reserves.

Theorem 12. *If a mechanism M is μ-revenue covered with reserves $\mathbf{r} = (r_1, \ldots, r_n)$, then the welfare of M is a $(1+\mu)e/(e-1)$-approximation to the welfare of the optimal mechanism which only serves agents with value $v_i \geq r_i$.*

Proof. Let $\mathbf{x}^*$ be the welfare-optimal allocation rule, and consider some value profile $\mathbf{v}$. Lemma 8 with $x' = x_i^*(\mathbf{v})$ yields that for each $\mathbf{x}$ and value v_i,

$$v_i x_i(v_i) + T_i[x_i(v_i), x_i^*(\mathbf{v})] \geq \tfrac{e-1}{e} v_i x_i^*(\mathbf{v}).$$

As all bidders in BNE play participatory actions, summing over all agents with values above reserves and using revenue covering yields:

$$\sum\nolimits_{i:v_i>r_i} v_i x_i(v_i) + \mu \textsc{Rev}(M) \geq \tfrac{e-1}{e} \sum\nolimits_{i:v_i>r_i} v_i x_i^*(\mathbf{v}).$$

Taking expectation with respect to $\mathbf{v}$, noting that $\mathbf{x}^*$ only serves agents with $v_i \geq r_i$ and agents have value r_i with probability 0, we get $\textsc{Welfare}(M) + \mu \textsc{Rev}(M) \geq \tfrac{e-1}{e} \textsc{Welfare}(\textsc{Opt})$; hence $(1+\mu)\tfrac{e}{e-1}\textsc{Rev}(M) \geq \tfrac{e}{e-1}\textsc{Welfare}(\textsc{Opt})$. □

5 Revenue Approximation

Recall that by Myerson's characterization of Bayes-Nash equilibrium (Theorem 1), the expected revenue can be viewed as the expected virtual welfare of agents served (Lemma 2). We will consider the task of approximating the revenue of the optimal auction in two parts: showing that the virtual welfare from positive virtual-valued agents approximates the optimal revenue, and demonstrating a few methods to ensure that the virtual welfare from agents with negative virtual values does not hurt revenue too much.

5.1 Positive Virtual Value Approximation

In Theorem 6 of Section 3, we showed that the first-price auction with monopoly reserves had approximately optimal revenue, via a reduction to the welfare approximation result. We show in this section that the same approach suffices to show that for any μ-revenue covered mechanism, the revenue accounted for by positive virtual valued agents approximates the optimal revenue.

Definition 13. *Let the positive and negative virtual values for an agent be $\phi_i^+(v_i) = \max(\phi_i(v_i), 0)$ and $\phi_i^-(v_i) = \min(\phi_i(v_i), 0)$ respectively. Define the positive and negative virtual welfare of a mechanism to be $\textsc{Rev}^+(M) = \sum_i \mathbf{E}_{v_i}[\phi_i^+(v_i)x_i(v_i)]$ and $\textsc{Rev}^-(M) = -\sum_i \mathbf{E}_{v_i}[\phi_i^-(v_i)x_i(v_i)]$.*

By Theorem 1, $\textsc{Rev}(M) = \textsc{Rev}^+(M) - \textsc{Rev}^-(M)$. Our primary result in this section is that $\textsc{Rev}^+(M)$ is a constant approximation to the revenue of the optimal mechanism if M is μ-revenue covered. Thus, bounding the loss from $\textsc{Rev}^-$ as a fraction of $\textsc{Rev}^+$ is sufficient to show approximately optimal revenue.

Theorem 14. *In any BNE of a μ-revenue covered mechanism M with single-dimensional agents, the positive virtual welfare $\textsc{Rev}^+(M)$ is a $(\mu+1)\tfrac{e}{e-1}$ approximation to the revenue of the optimal mechanism. More precisely, $\textsc{Rev}^+(M) + \mu\textsc{Rev}(M) \geq \tfrac{e-1}{e}\textsc{Rev}(\textsc{Opt})$.*

Recall that the approximation bound for μ-revenue covered auctions (Theorem 10) relied on showing that the surplus from any agent in any alternate allocation was approximated by that player's contribution to BNE surplus and a fraction of the additional threshold.

We begin by showing virtual-value covering, an analogue of value covering for virtual welfare, holds in BNE directly via a reduction to value-covering (Lemma 8).

Lemma 15 (Virtual-Value Covering)**.** *Consider a mechanism M in BNE and an agent i with value v_i. For any $x' \in [0,1]$,*

$$\phi_i^+(v_i)x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e}\phi_i^+(v_i)x'. \tag{15}$$

The proofs of Theorem 14 and Lemma 15 follow precisely as in Lemma 7 via a reduction to value-covering (Lemma 8), so the details are omitted.

Now that the positive virtual welfare of a mechanism approximates the optimal, the only thing left is to bound the loss due to serving bidders with negative virtual values. The subsequent sections discuss methods for mitigating the virtual welfare lost from serving negative virtual valued agents.

5.2 Reserve Prices

The standard approach to prevent service to agents with negative virtual values is to set reserves such that no negative virtual valued agent is served. As long as the virtual value $\phi_i(v_i)$ is non-decreasing in v_i — equivalently, the distribution is regular — setting monopoly reserves $\mathbf{r}^*$ s.t. $r_i^* = \phi_i^{-1}(0)$ in a first-price auction for every agent will eliminate all negative virtual valued agents. If a general mechanism serves no agent with $\phi_i(v_i) < 0$, then we will say it implements the reserves, and it too will approximate the revenue of the optimal mechanism if it is revenue covered above the reserves (recall Definition 11).

Definition 16. *A mechanism M implements reserves $\mathbf{r}$ if it serves no agents with values below their reserves in BNE, i.e. $x_i(v_i) = 0$ if $v_i < r_i$.*

Lemma 17. *Let M be a mechanism with agents having single-dimensional, regularly-distributed valuations. If M implements monopoly reserves $\mathbf{r}^*$ and is μ-revenue covered with monopoly reserves $\mathbf{r}^*$, then the revenue of M in any BNE is a $(\mu+1)e/(e-1)$ approximation to the revenue of the optimal mechanism.*

The proof is straightforward — as M serves no agent with $\phi_i(v_i) < 0$, $\textsc{Rev}^-(M) = 0$. By Theorem 14, M is then a $(\mu+1)e/(e-1)$ approximation to the revenue optimal mechanism. Thus if it is possible to add monopoly reserves to a mechanism, doing so gives approximately optimal revenue.

In a first price auction it is always feasible to implement reserves by explicitly adding the reserves, and in fact adding reserves preserves revenue covering (above the reserves):

Lemma 18. *If M is a μ-revenue covered, first-price auction, the mechanism $M^{\mathbf{r}}$ created by adding bid-space reserves $\mathbf{r} = (r_1, \ldots, r_n)$ to M is μ-revenue covered with reserves $\mathbf{r}$.*

Proof. Consider a strategy profile $\mathbf{s}$ from $M^{\mathbf{r}}$, with any bids from agents with values below their reserves mapped to 0. This strategy profile will have the same outcomes and revenues in M or $M^{\mathbf{r}}$, and for every agent i, every bid $b_i \geq r_i$ will result in the same allocation. Thus, for any participatory action profile a in $M^{\mathbf{r}}$, $T[\tilde{x}_i(a_i), x_i'] = T^{\mathbf{r}}[\tilde{x}_i(a_i), x_i']$. Since M is μ-revenue covered, it is also μ-revenue covered with reserves $\mathbf{r}$, and we have

$$\mu\textsc{Rev}(M^{\mathbf{r}}) = \mu\textsc{Rev}(M) \geq \sum\nolimits_{i:v_i>r_i} T_i[\tilde{x}_i(a_i), x_i'] = \sum\nolimits_{i:v_i>r_i} T_i^{\mathbf{r}}[\tilde{x}_i(a_i), x_i'].$$

Thus $M^{\mathbf{r}}$ is μ-revenue covered with reserves $\mathbf{r}$. ∎

In an all-pay auction however, we cannot reliably implement reserves. The willingness of a player to outbid an all-pay reserve depends on the allocation probability as well as the reserve, and as such there is no easy correspondence between all-pay and value space reserves. We now move to showing that adequate competition in the form of duplicate bidders is also sufficient to mitigate the impact of bidders with negative virtual values on revenue.

5.3 Duplicate bidders

Another approach to mitigating the impact of negative virtual-valued agents is to ensure each agent faces adequate competition. Bulow and Klemperer [1996] show that this intuition guarantees approximately optimal revenue in regular, symmetric, single-item settings.

We show the same intuition holds in μ-revenue covered mechanisms: if each bidder must compete with at least $k-1$ other bidders with values drawn from her same distribution and bidders play by identical strategies, revenue is approximately optimal compared to the revenue optimal mechanism (including the duplicate bidders). We say such an auction satisfies *k-duplicates*.

Lemma 19. *For any mechanism M with k-duplicates behaving by identical strategies and having values drawn from regular distributions, the virtual surplus lost due to serving agents with negative virtual values is at most $1/k$ the virtual surplus from positive virtual valued agents in any BNE, hence* $\textsc{Rev}(M) \geq \frac{k-1}{k}\textsc{Rev}^{+}(M)$.

The proof of Lemma 19 is included in Appendix C. Combining Lemma 19 with Theorem 14 thus ensures approximately optimal revenue:

Corollary 20. *For any mechanism M with k-duplicates behaving by identical strategies and having values drawn from regular distributions, the revenue is a $\left(\frac{k}{k-1}+\mu\right)\frac{e}{e-1}$-approximation to the revenue of the optimal mechanism.*

In both first price and all-pay auctions, duplicates will play by identical strategies and thus each will give approximately optimal revenue. Chawla and Hartline [2013] show that in a single-item setting, all bidders in a class that includes first-price and all-pay auctions (rank-and-bid based allocation rules, and bid-based payments) will behave symmetrically in BNE. If a mechanism has k-duplicates with such a payment rule, then for any group of duplicates, competing for allocation appears like a single-item auction, since at most one bidder of the group can be served. Thus, Theorem 3.1 of Chawla and Hartline [2013] will imply that agents in the same group behave by identical strategies:

Corollary 21 (of Theorem 3.1, Chawla and Hartline [2013])**.** *In any BNE of an auction with k-duplicates, rank-and-bid based allocation and bid-based payment, for any group B_j of agents, all agents in the group play by identical strategies everywhere except on a measure zero set of values.*

Thus Corollary 20 will hold for μ-revenue covered first price and all-pay auctions.

6 Revenue Covering

We have now shown that revenue-covering is a sufficient condition for a number of welfare and revenue approximation results:

- If M is μ-revenue covered, the welfare of M is a $(1+\mu)\frac{e}{e-1}-$approximation to the optimal welfare. (Theorem 10 in Section 4.2)
- If M implements monopoly reserves $\mathbf{r}^*$ and is μ-revenue covered with $\mathbf{r}^*$, the revenue of M is a $(1+\mu)\frac{e}{e-1}-$approximation to the optimal revenue. (Lemma 17 in Section 5.2)
- If M is μ-revenue covered and implements $k-$duplicates, the revenue of M is a $(\frac{k}{k-1}+\mu)\frac{e}{e-1}$-approximation to the optimal revenue. (Corollary 20 in Section 5.3)

In this section, we show that many auctions beyond the single-item, first-price auction are revenue covered, and thus get new welfare and revenue approximation results for each.

6.1 First Price Matroid Auctions

In our discussion of the single-item case (Section 3), we saw that when an agent has trouble getting allocated in a first price auction (that is, the cumulative additional threshold T_i is high), it is because other agents submit high bids. These competing bids translate into revenue, implying that the first-price auction is 1-revenue covered. With one extra step, this reasoning extends to first-price auctions where the feasible allocations form a matroid. An agent's threshold bid does not precisely correspond to a competing bid, but matroid properties provide a sufficiently close analog, implying revenue covering. Combining with Theorem 10 and Lemma 17 respectively imply welfare and revenue approximations of $2e/(e-1)$ with reserves. If the auction implements k-duplicates, which happens in the single-item setting when at least k bidders have values drawn from each distribution, then by Corollary 20 it is a $\frac{k}{k-1}\frac{2e}{e-1}$ approximation to the revenue optimal mechanism.

Lemma 22. *The first-price auction is 1-revenue covered in any matroid feasibility environment.*

The proof is included in Appendix D.

6.2 Position Auctions

In first-price position auctions (a.k.a., the generalized first-price auction, GFP), arguments similar to those in the matroid case yield analogous welfare and revenue guarantees.

Formally, a position auction is an auction in which agents can be allocated one of m positions; each of which is valued by an agent at $\alpha_j v_i$. In advertising auctions, these are slots on a webpage to fill where lower slots receive fewer clicks. The positions are ordered such that $\{\alpha_j\}$ is decreasing in j (hence slot 1 is best).

In GFP, agents submit bids b_i, and positions are allocated in order of bid. Each agent pays their bid scaled by the quality of the slot: $\alpha_j b_i$. Equivalently, they pay their bid when they are served, which occurs with probability α_j for position j.

While the correspondence between bids and threshold bids is not as immediate in GFP as in the single-item, first-price auction, GFP satisfies a notion of pointwise revenue covering, or revenue covering when other players are playing fixed actions that we use to show that it satisfies general revenue covering. The proof is included in D. By Lemma 18, adding reserves preserves revenue covering above the reserves and so Theorem 10 and Lemma 17 respectively imply welfare and revenue approximations of $2e/(e-1)$ with reserves.

Theorem 23. *The generalized first price (GFP) auction is* 1*-revenue covered.*

6.3 All-Pay Auctions

Revenue-covering is not limited to first-price semantics: the all-pay is also revenue covered, with $\mu = 2$. The relationship between revenue and thresholds in an all-pay auction is more indirect than for the first price auction. In a first-price auction, if a bidder must bid \$10 to attain a 80% probability of winning, it must be because the highest bid from other bidders is greater than \$10 80% of the time. In an all-pay auction however, the bidder would have to always pay \$10 instead. Recall that the equivalent bid of an action is the price per unit of allocation, so the equivalent bid of the \$10 all-pay bid giving allocation of 80% is \$10/80% = \$12.50 which is more than the competing bids that induced the threshold.

As the allocation probability from an all-pay bid approaches 1, the equivalent bid approaches the all-pay bid, and so such bids are almost revenue-covered. For small probabilities of allocation however, the all-pay bid and equivalent bid may be very far apart. The proof uses the monotonicity of the allocation rule to show that the threshold from bids with high probabilities of allocation outweighs the threshold from bids with low probabilities of allocation, giving a bound for the cumulative additional threshold of at most twice the revenue of the auction.

Combining with Theorem 10 gives a welfare bound of $3e/(e-1)$, weaker than the bound of 2 in Syrgkanis and Tardos [2013]. For revenue, it is not feasible to add reserves to an all-pay auction, but the all-pay auction with at least k-bidders from each distribution satisfies k-duplicates, and thus by Corollary 20 the revenue of the all-pay auction with at least 2 bidders with values from each distribution is a $4e/(e-1)$-approximation to the revenue of the optimal auction.

Lemma 24. *The all-pay matroid auction is 2-revenue covered.*

We include the proof here for the single-item case; the generalization to matroids is included in Appendix D.

Proof (single-item). We first translate revenue to threshold bids. In expectation, these thresholds 2-approximate the equivalent threshold bids. Combining the two arguments yields the result.

Let $\mathbf{s}$ be an arbitrary strategy profile, and $\mathbf{x}'$ an alternate allocation. First, for any bidder i, let $\tau_i^b(\mathbf{v}_{-i})$ be the threshold (all-pay) bid for i in realized value profile $\mathbf{v}_{-i}$ under $\mathbf{s}$. Since the threshold bid corresponds to some other agent's bid, and agents pay their bids regardless of allocation, $\textsc{Rev}(M) \geq \mathbf{E}_{\mathbf{v}_{-i}}[\tau_i^b(\mathbf{v}_{-i})] \geq \mathbf{E}_{\mathbf{v}_{-i}}[\tau_i^b(\mathbf{v}_{-i})]x_i'$.

To relate threshold bids to equivalent thresholds, let $a_i(z)$ be the z-quantile of i's competing bids. That is, $a_i(z) = \arg\min_{a_i} \tilde{p}_i(a_i)$ subject to $\tilde{x}_i(a_i) \geq z$. By the definition of τ_i,

$$\frac{\tilde{p}_i(a_i(z))}{\tilde{x}_i(a_i(z))} \geq \tau_i(z).$$

Rearranging and noting that in an all-pay auction, $\tilde{p}_i(a_i(z)) = a_i(z)$, we obtain

$$a_i(z) \geq \tau_i(z)\tilde{x}_i(a_i(z)) \geq \tau_i(z)z. \tag{16}$$

This yields the following sequence of inequalities:

$$\mathbf{E}_{\mathbf{v}_{-i}}\left[\tau_i^b(\mathbf{v}_{-i})\right] = \int_0^1 a_i(z)\,dz \geq \int_0^1 \tau_i(z)z\,dz \geq \frac{1}{2}\int_0^1 \tau_i(z)\,dz = T_i[0,1], \tag{17}$$

where the first equality follows from noting that expected value can be computed by integrating over quantiles, the first inequality from equation (16), and the second inequality from the fact that τ_i is an increasing function and Chebyshev's sum inequality. Finally, since $\mathbf{x}'$ is feasible, $\sum_i x'_i \leq 1$. We can combine this with (17) to get

$$2\textsc{Rev}(M) \geq \sum\nolimits_i T_i[0,1]x'_i. \tag{18}$$

By the convexity of T_i, $T_i[0,1]x'_i \geq T_i[0,x'_i]$. Since $T_i[0,x'_i] \geq T_i[\tilde{x}_i(a_i), x'_i]$, we conclude that the all-pay auction is 2-revenue covered. □

6.4 The Second-Price Auction

Not all mechanisms are revenue covered. One such mechanism that lacks a direct relationship between thresholds and revenues is the second-price auction. In the second-price auction, agents submit sealed bids, the highest bidder wins and is charged the second-highest bid. Consider a two-agent setting where bidders have deterministic values $v_1 = 1$ and $v_2 = \epsilon$. Assume agent 1 bids 1 and agent 2 bids ϵ. The revenue is ϵ, but $T_2[0,1]$ is 1, so the second-price auction cannot be revenue covered.

7 Extension: Simultaneous Composition

In this section we prove that if a set of mechanisms satisfy revenue covering when operated in isolation, then they satisfy revenue covering when many instances of the mechanisms are simultaneously being run if agents are unit-demand and single-valued across outcomes of the mechanisms. We formally define simultaneous composition for single-dimensional agents as follows:

Definition 25. *Let mechanisms $M_1, \ldots, M_m$ have allocation and payment rules $(\mathbf{x}^j, \mathbf{p}^j)$ and individual action spaces spaces $A_i^1, \ldots, A_i^m$ for each agent i. The* simultaneous composition *of $M_1, \ldots, M_m$ is defined to have:*

- *Action space $\prod_j A_i^j$ for each agent. That is, each agent participates in the global mechanism by participating in each composed mechanism individually.*
- *Allocation rule $\tilde{x}_i(\mathbf{a}) = \left[\tilde{x}_i^1(\mathbf{a}^1), \ldots, \tilde{x}_i^m(\mathbf{a}^m)\right]$. In other words, the mechanism gives each agent their allocated bundle from each mechanism.*
- *Payment rule $\tilde{p}_i(\mathbf{a}) = \sum_j \tilde{p}_i^j(\mathbf{a}^j)$. That is, agents make payments to every composed mechanism.*

We assume agent utilities are *unit demand* and *single-valued* over the outcomes of the mechanisms. Agent utilities are then of the form $v_i \cdot (\max_{j \in \mathcal{S}_i} \tilde{x}_i^j(\mathbf{a})) - \tilde{p}_i(\mathbf{a})$. The *induced single-dimensional allocation rule* is $\hat{x}_i(\mathbf{a}) = \max_{j \in \mathcal{S}_i} \tilde{x}_i^j(\mathbf{a})$; it serves the same role in our framework as the allocation rule of a single, uncomposed mechanism. We define τ_i, T_i, and revenue covering with respect to $\hat{x}_i$ for mechanisms defined as in previous sections.

The main theorem of this section is that the simultaneous composition of a number of mechanisms that satisfy the same notion of revenue covering (no reserves, or above the same individual reserves) remains revenue covered in the same manner.

Lemma 26. *Let M be the simultaneous composition of μ-revenue covered (with reserves $\mathbf{r}$) mechanisms $M_1, \dots M_m$ with unit-demand, single-valued agents; then M is μ-revenue covered (with reserves $\mathbf{r}$).*

The proof is included in Appendix E. Combining Lemma 26 with Theorem 10 and Lemma 17 immediate yields welfare and revenue results, for instance:

Corollary 27. *Let M be the simultaneous composition of k first-price auctions with monopoly reserves $\mathbf{r}^*$, matroid feasibility constraints, and unit-demand, single-valued agents with regularly distributed valuations. Then, the revenue of M is a $\frac{2e}{e-1}-$approximation to the revenue of the optimal global mechanism.*

8 Conclusion

We have shown a framework for proving price of anarchy results for welfare and revenue in Bayes-Nash Equilibrium. This framework enabled us to prove both welfare and new revenue approximation results for non-truthful auctions in asymmetric settings, including first price and all-pay auctions.

We split this framework in two distinct parts that isolate the analysis of Bayes-Nash Equilibrium from the analysis of the specific mechanism. The first part, value covering, depends only on Bayes-Nash Equilibrium and relates an agents surplus and expected price for additional allocation with her optimal surplus. The second, revenue-covering, depends only on properties of a mechanism over individually rational strategy profiles and feasible allocations. This is especially helpful when equilibria are hard to characterize or understand analytically, as is the case with the first-price auction in asymmetric settings. We expect this framework will aid broadly in understanding properties of equilibria in auctions well beyond the confines of symmetric settings.

We invoked the characterization of Bayes-Nash Equilibrium in a few very specific places in our proofs. For value covering and virtual value covering, it is only important that an agent be best responding to the expected actions of other bidders. For the revenue approximation results, we do rely on the characterization of equilibrium by Myerson [1981] to account for revenue via virtual values. This is the crucial part that allows us to relate the allocation a bidder receives to their contribution to revenue. Extensions beyond single-parameter, risk-neutral, private-valued agents will be challenging without a virtual-value equivalent.

References

Ausubel, L. M. and Milgrom, P. (2006). The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1.* MIT Press.

Bulow, J. and Klemperer, P. (1996). Auctions Versus Negotiations. *The American Economic Review*, 86(1):180–194.

Chawla, S. and Hartline, J. (2013). Auctions with Unique Equilibria. In *ACM Conference on Electronic Commerce*.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.

Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.

Jiang, A. X. and Leyton-Brown, K. (2010). Bayesian action-graph games. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A., editors, *Advances in Neural Information Processing Systems 23*, pages 991–999. Curran Associates, Inc.

Kaplan, T. R. and Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302.

Kirkegaard, R. (2009). Asymmetric first price auctions. *Journal of Economic Theory*, 144(4):1617–1635.

Kirkegaard, R. (2012). A Mechanism Design Approach to Ranking Asymmetric Auctions. *Econometrica*, 80(5):2349–2364.

Krishna, V. (2009). *Auction Theory*. Academic Press/Elsevier.

Lebrun, B. (2006). Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, 55(1):131–151.

Maskin, E. and Riley, J. (2003). Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395 – 409.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73.

Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *STOC*.

Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *ACM Conference on Electronic Commerce*, pages 844–861.

Syrgkanis, V. and Tardos, E. (2013). Composable and Efficient Mechanisms. In *STOC*.

Talwar, K. (2003). The price of truth: Frugality in truthful mechanisms. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science*, pages 608–619.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37.

A Examples of FPA Equilibria

A.1 Single-item FPA PoA ≥ 1.15

Consider a setting with one high bidder with a fixed value of 10, and n small bidders with values drawn from some distribution with value always less than 10. The welfare-optimal allocation always serves the high bidder. We parameterize the expected utility of the high bidder as u_H. Assume the low bidders will bid such that the highest of their bids is distributed according to the CDF $F_L(b) = u_H/(10-b)$, with a point mass of probability $u_H/10$ at 0. With this distribution, player H achieves utility u_h for any bid in the range $[0, 10-u_h]$.

The high player plays a mixed strategy according to the bid CDF $B_H(b) = \sqrt{b/(10-u_H)}$. The competing bid CDF for each low bidder is $F_c(b) = B_H(b) \cdot B_L(b)^{(n-1/n)}$.

With $u_H = 5.7$, solving for the first order conditions in the first price auction tells us that for any low player bidding b, $v = b + F_c(b)/F_c'(b)$; solved numerically it is approximately $v(b) = \frac{15b-0.5b^2}{5+0.5b}$. Solving numerically gives welfare of 8.69; thus the price of anarchy for welfare is approximately 1.15.

This example is almost tight against the expected cumulative threshold lowerbound $\underline{T}$ used in the proof of the value covering lemma (Lemma 8). However, the $\frac{e}{e-1}$ price of anarchy proof ignores the bid from the agent allocated in the optimal allocation and the utility of the agents allocated in FPA but not Opt. Both of these quantities are non-zero, which leads to the 1.15 figure being reasonably far from the $\frac{e}{e-1}$. Bounding these quantities is a likely required step for improving the $\frac{e}{e-1}$ bound for single-item settings.

B Framework Proofs

Lemma 8 (Restatement). *Consider a mechanism M in BNE with induced allocation and payment rules $(\mathbf{x}, \mathbf{p})$, and an agent i with value v_i. For any $x' \in [0,1]$,*

$$v_i x_i(v_i) + T_i[x_i(v_i), x'] \geq \tfrac{e-1}{e} v_i x'. \tag{13}$$

Proof of Lemma 8. If $x_i(v_i) > x'$, $T_i[x_i(v_i), x'] = 0$ and the result follows. Otherwise, note that by the definition of BNE, i chooses an action which maximizes utility. It follows that

$$u_i(v_i) \geq v_i x_i(\alpha_i(z)) - p_i(\alpha_i(z)) = \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) x_i(\alpha_i(z)) \geq \left(v_i - \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))}\right) z. \tag{19}$$

Rearranging (19) yields

$$v_i - \frac{u_i(v_i)}{z} \leq \frac{p_i(\alpha_i(z))}{x_i(\alpha_i(z))} = \tau_i(z). \tag{20}$$

This bound is meaningful as long as $v_i - \frac{u_i(v_i)}{z} \geq 0$, or alternatively $z \geq u_i(v_i)/v_i$. It follows that

$$\begin{aligned} v_i x_i(v_i) + T_i[x_i(v_i), x'] &\geq v_i x_i(v_i) + \int_{x_i(v_i)}^{x'} \max\left(0, v_i - \frac{u_i(v_i)}{z}\right) dz \\ &\geq v_i x_i(v_i) + \int_{u_i(v_i)/v_i}^{x'} v_i - \frac{u_i(v_i)}{z}\, dz \\ &= v_i x' + u_i(v_i) \ln \frac{u_i(v_i)}{x' v_i}. \end{aligned} \tag{21}$$

Holding $x'v_i$ fixed and minimizing this quantity as a function of $u_i(v_i)$ yields a minimum at $u_i(v_i) = \frac{x'v_i}{e}$, and at that point assumes value $(1-1/e)x'v_i$. This is precisely the righthand side of (13), implying the lemma. □

C Revenue Extension Proofs

C.1 Duplicates Environment

We first formally define the k-duplicates environment.

Definition 28 (k-duplicates)**.** *An auction has* k-duplicates *if there is a partition of bidders into groups* $\{B_1, \ldots, B_p\}$ *each of size* $n_j \geq k$ *such that:*

- *the value of every bidder in each group is drawn from the same distribution* F_j*,*
- *at most one bidder from any group can feasibly be served, and*
- *each bidder in a group is treated identically by the auction.*

This is the same as the duplicates environment of Hartline and Roughgarden [2009]; however, our approximation results will hold with respect to the optimal mechanism in the duplicates environment, which is always better than the optimal auction in the same environment without duplicates. If duplicates play by identical strategies in BNE, this guarantees that a bidder can only be served if she has the highest value among her duplicates.

We assume for this proof that the allocation rule is pointwise monotonic for any bidder. This is stronger than the monotonicity assumption of the BNE characterization, but is satisfied in any bid-based mechanism (first-price, all-pay, etc.) with monotonicity of bids. A slightly more intricate analysis of the revenue curves of the duplicates as a group extends the argument here to the weaker monotonicity assumption.

We first relate the revenue from each group of bidders to the revenue from a symmetric second price auction with reserves among only the bidders within the group of duplicates, allowing us to use the symmetric auction approximation results of Bulow and Klemperer [1996].

Let $\text{SPA}_R(B)$ be a second price auction run among agents in group B with a random reserve drawn according to R.

Lemma 29. *There exist reserve value distributions* $R_1, R_2 \ldots R_p$ *such that in any mechanism* M *with* k*-duplicates,*

$$\textsc{Rev}(M) = \sum\nolimits_j \textsc{Rev}(\text{SPA}_{R_j}(B_j)), \tag{22}$$

$$\textsc{Rev}^+(M) = \sum\nolimits_j \textsc{Rev}^+(\text{SPA}_{R_j}(B_j)). \tag{23}$$

Proof. By the pointwise monotonicity of the allocation rule, fixing the values and actions of bidders outside a group j results in a threshold value for the top ranking member of a group. Let the distribution of such thresholds be R_j; then a second price auction among the group members with reserve drawn precisely from R_j will induce exactly the same allocation rule for all members of the group. By revenue equivalence (Part 3 of Theorem 1), the revenue from the group in mechanism M will be the same as $\textsc{Rev}(\text{SPA}_{R_j}(B_j))$. The same argument holds for $\textsc{Rev}^+(M)$ and $\textsc{Rev}^+(\text{SPA}_{R_j}(B_j))$. □

Lemma 19 (Restatement). *In any BNE of a mechanism* M *with* k*-duplicates behaving by identical strategies and having values drawn from regular distributions, the virtual surplus lost due to serving agents with negative virtual values is at most* $1/k$ *the virtual surplus from positive virtual valued agents.*

Proof of Lemma 19. A second-price auction within a group is now a symmetric setting, and thus we can now use the work of Bulow and Klemperer [1996] to relate (22) and (23). By Bulow and Klemperer [1996], if $k \geq 2$, $\textsc{Rev}(\text{SPA}_{R_j}(B_j)) \geq \frac{k-1}{k}\textsc{Rev}^+(\text{SPA}_{R_j}(B_j))$ and hence:

$$\begin{aligned}\textsc{Rev}(M) &= \sum\nolimits_j \textsc{Rev}(\text{SPA}_{R_j}(B_j)) \\ &\geq \sum\nolimits_j \frac{k-1}{k}\textsc{Rev}^+(\text{SPA}_{R_j}(B_j)) \\ &= \frac{k-1}{k}\textsc{Rev}^+(M).\end{aligned}$$

Thus $\textsc{Rev}^- \leq \frac{1}{k}\textsc{Rev}^+(M)$, exactly our desired result. □

D Revenue Covering Proofs

D.1 First Price Matroid Auctions

To prove revenue covering for matroids, we first make note of a property relating threshold bids to bids in bid-based matroid auctions, related to a property shown in Talwar [2003] for VCG in matroid environments. For first-price auctions, threshold bids are the same as equivalent threshold bids and so this will lead to revenue covering directly; for all-pay auctions, bid thresholds and equivalent bid thresholds can be different and depend on the probability of allocation, so the proof is not as immediate.

Lemma 30. *Let M be a bid-based mechanism with a matroid feasibility constraint that allocates so as to maximize the sum of bids from served agents. Let $\tau_i^b(\mathbf{v}_{-i})$ be the threshold bid for bidder i given actions $\mathbf{s}_{-i}(\mathbf{v})$. Then for any strategy profile $\mathbf{s}$, value profile $\mathbf{v}$ and feasible allocation $\mathbf{x}'$,*

$$\sum\nolimits_i s_i(v_i)x_i(\mathbf{v}) \geq \sum\nolimits_i \tau_i^b(\mathbf{v}_{-i})x_i'. \tag{24}$$

The proof is based on the following property of matroids:

Lemma 31 (Replacement Property)**.** *Let S_1 and S_2 be independent sets of size k in a matroid $\mathcal{M}$. Then there is a bijective function $f : S_2 \setminus S_1 \to S_1 \setminus S_2$ such that, for every $i \in S_2 \setminus S_1$, the set $(S_1 \setminus \{f(i)\}) \cup \{i\}$ is independent in $\mathcal{M}$.*

Proof of Lemma 30. Because subsets of feasible allocations are feasible, threshold bids are nonnegative, so we only need consider allocations $\mathbf{x}'$ which are bases. Let S and S' be sets served by $\mathbf{x}$ and $\mathbf{x}'$, respectively. Since bids are nonnegative, it follows that S and S' are the same size. By Lemma 31, there exists a bijection f from $S' \setminus S$ to $S \setminus S'$ with the replacement property in the lemma.

For each $i \in S' \setminus S$, $s_{f(i)}(\mathbf{v}_{f(i)}) \geq \tau_i^b(\mathbf{v}_{-i})$, as if i bids above $s_{f(i)}$, then $(S \setminus \{f(i)\}) \cup \{i\}$ would be optimal and therefore i would be allocated in BNE. For each $i \in S' \cap S$, i was served in $x(\mathbf{v})$, it must be that $s_i(v_i) \geq \tau_i^b(\mathbf{v}_{-i})$. The result follows by summing over i. □

With a relationship between threshold bids and revenue established, it remains to connect the threshold bids to T_i. We already saw in the proof of Lemma 5 that this is simple. With first-price semantics, $\tau_i(z)$ is simply the z-quantile of threshold bids. It follows that $T_i[0,1]$ is i's expected threshold bid. Using this relationship, we get:

Lemma 22 (Restatement). *The first-price auction is 1-revenue covered in any matroid feasibility environment.*

Proof. Consider some alternate allocation $\mathbf{x}'$ and action profile $\mathbf{a}$. By the mechanism's payment scheme and Lemma 30,

$$\textsc{Rev}(M) = \mathbf{E}_{\mathbf{v}}\left[\sum_i s_i(v_i)x_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i \tau_i^b(\mathbf{v}_{-i})x_i'\right] = \sum_i \mathbf{E}_{\mathbf{v}}\left[\tau_i^b(\mathbf{v}_{-i})\right] x_i'.$$

It follows that $\mathbf{E}_{\mathbf{v}}[\tau_i^b(\mathbf{v}_{-i})] = T_i[0,1]$. Using this fact, we get $\textsc{Rev}(M) \geq \sum_i T_i[0,1]x_i'$. Finally, the convexity of T_i yields that $\sum_i T_i[0,1]x_i' \geq \sum_i T_i[0,x_i']$. Finally, noting that $\sum_i T_i[0,x_i'] \geq \sum_i T_i[\tilde{x}_i(a_i),x_i']$ proves the lemma. □

Combining Lemma 22 with Theorem 10 and yields

Corollary 32. *For the first price matroid auction with arbitrary reserves, the welfare of any BNE is a $2e/(e-1)$-approximation to that of any other mechanism with those same reserves.*

Moreover, using Lemma 17 and Lemma 19, we get

Corollary 33. *For the first price matroid auction with monopoly reserves and regular bidders, the revenue of any BNE is a $2e/(e-1)$-approximation to that of any other mechanism.*

Corollary 34. *For the first price matroid auction with regular bidders and at least 2 duplicates, the revenue of any BNE is a $3e/(e-1)$-approximation to that of any other mechanism.*

D.2 All-Pay Matroid Auction

Lemma 24 (Restatement). *The all-pay matroid auction is 2-revenue covered.*

Proof. The proof for matroid environments differs from the single-item proof only in the use of Lemma 30 to relate revenue to threshold bids.

Revenue to Threshold Bids. Let $\mathbf{x}'$ be a feasible allocation and $\mathbf{a}$ be an action profile. By the payment semantics of the mechanism,

$$\textsc{Rev}(M) = \mathbf{E}_{\mathbf{v}}\left[\sum_i s_i(v_i)\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i s_i(v_i)x_i(\mathbf{v})\right].$$

Now let $\tau_i^b(\mathbf{v}_{-i})$ be the threshold bid for i in realized value profile $\mathbf{v}_{-i}$ under strategy profile $\mathbf{s}$ (without index i). Because the served agents are the basis which maximizes the sum of bids, Lemma 30 implies that

$$\mathbf{E}_{\mathbf{v}}\left[\sum_i s_i(v_i)x_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i \tau_i^b(\mathbf{v}_{-i})x_i'\right] = \sum_i \mathbf{E}_{\mathbf{v}}\left[\tau_i^b(\mathbf{v}_{-i})\right] x_i'. \tag{25}$$

The rest of the proof proceeds exactly the same as the proof for the single-item case: translating between threshold bids and equivalent threshold bids (losing the factor of 2), and finally summing over all agents to achieve the desired results. □

D.3 GFP

In deterministic mechanisms, we used the pointwise equivalent bid threshold for allocation $\tau_i(\mathbf{v}_{-i})$, or the required bid to be allocated when other agents have values $\mathbf{v}_{-i}$. In a randomized mechanism like a position auction, fixing the actions of other results not in a single threshold but a number of thresholds — the actions of others induce an allocation rule that in the case of position auctions, is piecewise constant (a "stair" function).

We will make use of the threshold that is induced by this action profile in proving that GFP is revenue-covered. Let $\tau_i^{\mathbf{a}_{-i}}(z) = \beta_i(\alpha_i(z, \mathbf{a}_{-i}), \mathbf{a}_{-i})$ be the smallest equivalent bid of an action for bidder i which achieves at least allocation of z when other bidders play actions $\mathbf{a}_{-i}$. Let $T_i^{\mathbf{a}}[x_i(\mathbf{a}), x'] = \int_{x_i(\mathbf{a})}^{x'} \tau_i^{\mathbf{a}}(z)\, dz$ denote the expected additional threshold for agent i when other bidders play $\mathbf{a}_{-i}$.

To prove GFP is revenue covered for all strategy profiles, we will show first that GFP satisfies a pointwise variant of revenue covering; then, that pointwise revenue covering implies revenue covering.

Definition 35. *A mechanism M is pointwise μ-revenue covered if for any actions $\mathbf{a}$, participatory actions $\mathbf{a}'$ and alternate feasible allocation x',*

$$\mu REV(M(\mathbf{a})) \geq \sum_i T_i^{\mathbf{a}_{-i}}[\tilde{x}_i(a_i', \mathbf{a}_{-i}), x_i']. \tag{26}$$

Lemma 36. *GFP is pointwise 1-revenue covered.*

Proof. Consider the bid-based allocation rule of an agent in GFP, $\tilde{x}_i(b_i, \mathbf{b}_{-i})$. For any bid b_i, $\tilde{x}_i(b_i, \mathbf{b}_{-i})$ is the position weight of the best slot such that the current resident of the slot is bidding less than b_i. So, $\tilde{x}_i(b_i)$ will be a stair function, with a stair corresponding to each position. The area above the curve between allocation probabilities $\tilde{x}_i(b_i, \mathbf{b}_{-i})$ and x_i', $T_i^{\mathbf{a}}[\tilde{x}_i(b_i, \mathbf{b}_{-i}), x']$, is a lower bound on the actual payment made by the bidder in the slot, because it assumes that the current winner is paying his bid only for the extra marginal clicks, not the clicks across all the slots. Denote by b^j the winning bid for each position j; then

$$T_i^{\mathbf{a}}[\tilde{x}_i(b_i, \mathbf{b}_{-i}), \alpha_j] = \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b^j. \tag{27}$$

The revenue in GFP given a set of bids is $\sum_j \alpha_j b^j$. For any slot j, the threshold amount for the bidder allocated j in the alternate allocation is less than payment of the bidder who won the slot j: $\alpha_j b^j \geq \sum_{i=j}^{m} (\alpha_i - \alpha_{i+1}) b^j$. Summing over all bidders gives that $REV(\mathbf{a}) \geq \sum_i T_i^{\mathbf{a}}[\tilde{x}_i(b_i, \mathbf{b}_{-i}), x_i']$, our desired result. □

We now prove (general) revenue covering for the generalized first price auction:

Theorem 23 (Restatement). *GFP is 1-revenue covered.*

Proof. Let $\mathbf{x}'$ be a feasible allocation. Applying pointwise revenue covering with every set of realized actions $\mathbf{s}(\mathbf{v})$, and with an alternate action profile of always bidding 0, gives

$$REV(GFP) \geq \mathbf{E}_{\mathbf{v}} \left[\sum_i T_i^{\mathbf{s}(\mathbf{v})}[0, x_i'] \right]. \tag{28}$$

If the expected pointwise threshold on the right side of (28) was the same as the cumulative additional threshold of the general mechanism, our proof would be complete. Unfortunately it is not, and instead gives the threshold if you were able to bid to get exactly the same slot for every set of realized bids from other agents — which is not generally a strategy that can be implemented by a bidder in GFP. However, even if agent i *could* react to the actions of all other agents, by the convexity of $T_i^{\mathbf{s}(\mathbf{v})}[0, x_i']$ in x_i', bidding $\tau_i(x_i')$ all the time is the strategy that attains allocation x_i' with the smallest expected threshold. Thus,

$$\mathbf{E}_{\mathbf{v}}\left[\sum_i T_i^{\mathbf{s}(\mathbf{v})}[0, x_i']\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_i T_i^{\mathbf{s}(\mathbf{v})}[0, \tilde{x}_i(\tau_i(x_i'))]\right]. \tag{29}$$

The right side of equation (29) is exactly the cumulative additional threshold $\sum_i T_i[\tilde{x}_i(\mathbf{s}(\mathbf{v})), x_i']$, and so combining with equation (28) gives our desired result, $REV(GFP) \geq \sum_i T_i[x_i(v_i), x_i']$. □

E Simultaneous Composition Proofs

We prove revenue covering holds under simultaneous composition by proving a stronger notion of revenue covering holds, revenue covering when bidders have a restricted set of actions they can take. We consider thresholds derived from participating in only one of the mechanisms being composed. These thresholds will retain the revenue covering properties of the original mechanisms and cause the global mechanism to inherit the property as well.

Given a strategy $\mathbf{s}$ in the composed mechanism, let $\mathbf{s}^j$ denote the strategy profile in mechanism j defined by the element of each agent's strategy profile corresponding to M_j. Given $\mathbf{s}^j$, define τ_i^j, and T_i^j to be the analogous values of τ_i, and T_i in M_j under $\mathbf{s}^j$. In the composed mechanism, let A_i^j be the set of actions comprised of an arbitrary action in mechanism j and withdrawing from all other mechanisms. Further let $A_i' = \bigcup_j A_i^j$, and $\mathbf{A}' = \prod_i A_i'$.

Action Restrictions Restricting the set of allowed actions by bidders only decreases their allocation for any equivalent bid and hence only increases their threshold. It follows that if a mechanism's revenue covers the cumulative threshold derived from a restricted action set, then it also covers those from the unrestricted action set. The following are analogous to the original definitions, but with a restricted action set:

Definition 37. *The equivalent threshold bid with respect to a restricted action set $A_i' \subseteq A_i$, denoted $\tau_i^{A_i'}(z)$, is defined as $\min_{a_i \in A_i'} \beta_i(a_i)$ subject to $\tilde{x}_i(a_i) \geq z$.*

Definition 38. *Given an action a_i, the expected equivalent threshold bid with respect to a restricted action set $A_i' \subseteq A_i$, denoted $T_i^{A_i'}[\tilde{x}_i(a_i), x_i']$, is defined as $\int_{\tilde{x}_i(a_i)}^{x'} \tau_i^{A_i'}(z)\, dz$.*

Definition 39. *A mechanism M is μ-revenue covered restricted to actions $A_1' \subseteq A_1, \ldots, A_n' \subseteq A_n$ (with product space $\mathbf{A}'$) if for every strategy profile $\mathbf{s}$, alternate feasible allocation $\mathbf{x}'$, and profile of general actions $\mathbf{a}$, $\mu \textsc{Rev}(M) \geq \sum_i T_i^{A_i'}[\tilde{x}_i(a_i), x_i']$.*

Note that restricted revenue covering is stronger than unrestricted revenue covering, and as such proving a mechanism is revenue covered for a given restriction is sufficient to show that it is revenue covered with no restriction. This follows from the fact that $\tau_i^{A_i'}(z)$ is the

objective value to the same minimization problem as $\tau_i(z)$, but on a smaller feasible region, so $\tau_i^{A'_i}(z) \geq \tau_i(z)$ for all i and z. Integrating, we see that $T_i^{A'_i}[\tilde{x}_i(a_i), x'_i] \geq T_i[\tilde{x}_i(a_i), x'_i]$ for all i and x', so $\textsc{Rev}(M) \geq \sum_i T_i^{A'_i}[\tilde{x}_i(a_i), x'_i] \geq \sum_i T_i[\tilde{x}_i(a_i), x'_i]$.

Additionally, restricted revenue covering combines with revenue covering for bidders with values above a set of reserves (Definition 11) in the logical manner, only summing over bidders with values above reserves.

Lemma 26 (Restatement). *Let M be the simultaneous composition of μ-revenue covered (with reserves $\mathbf{r}$) mechanisms $M_1, \ldots M_m$ with unit-demand, single-valued agents. Then M is μ-revenue covered (with reserves $\mathbf{r}$).*

Proof. We prove that M is μ-revenue covered with respect to $\mathbf{A}'$, the restriction to only participating in one mechanism at a time. We will assume for the proof that $\mathbf{r} = (0, \ldots, 0)$. We conclude with an explanation of the proof for the non-zero reserves case.

Let $\mathbf{x}'$ be a feasible induced allocation for the global mechanism. That is, we can construct a matching between agents and mechanisms such that for any j, there is a feasible allocation for M_j that allocates each i matched to j according to x'_i. Define $x'_{i,j}$ be x'_i if i and j are matched, and 0 otherwise. Note that for each agent i, $x'_{i,j} > 0$ for at most one j, with $x'_{i,j} = 0$ for all j if $x'_i = 0$. By downward closure, $\mathbf{x}'_{\cdot,j}$ is a feasible allocation for M_j.

Now by the definition of the composed mechanism, $\mu\textsc{Rev}(M) = \mu \sum_j \textsc{Rev}(M_j)$, where $\textsc{Rev}(M_j)$ is taken with respect to $\mathbf{s}^j$ for each j. Let $\mathbf{a}$ be an action profile in M, and let $\mathbf{a}^j$ be the corresponding vector of actions in mechanism j. Because each M_j is μ-revenue covered, it follows that $\mu \sum_j \textsc{Rev}(M_j) \geq \sum_j \sum_i T_i^j[\tilde{x}_i^j(a_i^j), x'_{i,j}]$. Moreover, for all j, $T_i^j[\tilde{x}_i^j(a_i^j), x'_{i,j}] = T_i^{A_i^j}[\tilde{x}_i^j(a_i^j), x'_{i,j}] \geq T_i^{A'_i}[\tilde{x}_i^j(a_i^j), x'_{i,j}]$, and by the definition of the induced single-dimensional allocation rule of the composed mechanism, $\hat{x}_i$, $T_i^{A'_i}[\tilde{x}_i^j(a_i^j), x'_{i,j}] \geq T_i^{A'_i}[\hat{x}_i(a_i), x'_{i,j}]$. But for each agent i, $x'_{i,j} > 0$ for at most one j, so $T_i^j[\tilde{x}_i^j(a_i^j), x'_{i,j}] > 0$ for at most one j as well, again with $T_i^j[\tilde{x}_i^j(a_i^j), x'_{i,j}] = 0$ for all j if $x'_i = 0$. The same also holds for $T_i^{A'_i}[\hat{x}_i(a_i), x'_{i,j}]$. It follows that $\sum_j \sum_i T_i^j[\tilde{x}_i^j(a_i^j), x'_{i,j}] \geq \sum_i T_i^{A'_i}[\hat{x}_i(a_i), x'_{i,j}]$, which implies the result. □

The proof with nonzero reserves differs slightly from the above in two ways. First, we consider the thresholds only of agents with values above their reserves. As above, the proof compares the thresholds of these agents in the global mechanism to those of the local mechanisms. Second, given a participatory action profile $\mathbf{a}$ for the global mechanism, we need only consider thresholds from mechanisms where a_i^j is also participatory. As reserves are the same across all mechanisms, the local cumulative thresholds above the allocation from a participatory action are larger than those for the global mechanism for the same allocation, giving the same result.